arXiv:1512.08077v1 [stat.ME] 26 Dec 2015

Model Prior Distribution for Variable Selection in Linear Regression Models

Cristiano Villa* and Jeong Eun Lee†

December 29, 2015

Abstract

In this work we discuss a novel model prior probability for variable selection in linear regression. The idea is to determine the prior mass in an objective sense, by considering the *worth* of each of the possible regression models, given the number of covariates under consideration. Through a simulation study, we show that the proposed prior outperforms the uniform prior and the Scott & Berger prior in a scenario of no prior knowledge about the size of the true regression models. We illustrate the use of the prior using two well-known data sets with, respectively, 15 and 4 covariates.

Keywords Bayesian variable selection, linear regression, loss functions, objective priors.

1 Introduction

In this paper, we propose a model prior probability distribution for variable selection problems in linear regression models. We focus on the general approach where, given d covariates, the task is to derive a posterior probability on the space of regression models. The prior we propose is based on losses and it is compared with both the uniform prior (which assigns equal prior mass to each model) and the Scott & Berger prior (Scott and Berger, 2010) through a simulation study. We show that, in absence of any prior information about the true regression model, the model prior we propose has an overall better performance than the other two. We also compare the model prior distributions on two well-known data sets: the US crime data (Vandaele, 1978) and the Hald data (Woods *et al.*, 1932).

Variable selection problems, in the Bayesian framework, are (and should be) in line with any other inferential procedure. That is, a posterior distribution for the space of models should be obtained in order to represent the (posterior) uncertainty about the true regression model (see Gelman *et al.* (2004)). There may be instances where the above is not appropriate, for example if there are models with a negligible posterior probability, in which case a subset of all the possible regression models can be considered. With a prior distribution on the space of models, representing the model uncertainty related to variable selection, one way to proceed is

*University of Kent, School of Mathematics, Statistics and Actuarial Sciences, Canterbury, UK E-mail: cv88@kent.ac.uk (corresponding author)

†Auckland University of Technology, School of Computing and Mathematical Sciences, Auckland, New Zealand E-mail: jeong.lee@aut.ac.nz

by using Bayesian model averaging (Hoeting *et al.*, 1999). When the model posterior distribution tends to be spread across many of the possible regression models, and when prediction is an important part of the statistical analysis, Raftery *et al.* (1997) show that Bayesian model averaging performs superiorly than choosing the regression model with the highest posterior probability. Also, although one may decide to explore a different route than model averaging, Barbieri and Berger (2004) show that the median probability model, under certain conditions, has a predictive power at least as good as the one of the highest probability model. The median probability model derives from an equivalent "answer" to the above problem obtained by estimating the marginal posterior probability that each variable has, independently from the others, of being included in the regression model.

An important component of the Bayesian variable selection approach is the definition of the prior for the regression coefficients, including the intercept, and the regression variance. In fact, we can comfortably say that the majority of literature related to variable selection is focussed on the identification of appropriate prior distributions for the model specific parameters. This is particularly the case when the approach has to be minimally informative; however, due to the large number of potential covariates examined in modern problems, a minimally informative prior is possibly the most wise (or the sole) solution as it would be not feasible to elicit priors for moderate to large models. Although the importance of parameter specific priors, in this paper we will focus on the model prior distribution only, referring to the specific literature on the subject. See, for example, Bayarri *et al.* (2012) and the references therein.

The prior we propose is based on the idea that, given a variable selection problem, where all the potential covariates have been identified, every possible linear regression model has a *worth* depending on the fact that that model has been chosen to be part of the problem (Villa and Walker, 2015). In other words, the choice of including a model in the model space conveys information, and we use that information to assign a prior probability to the model. It is important to highlight that the prior we propose is intended to be applied within the Bayesian framework, in the sense that the prior represents the initial uncertainty about the true model which is updated, by including the information coming from the data, to obtain a posterior distribution representing the uncertainty at the end of the process (Gelman *et al.*, 2004). The obtained prior depends on the model size, that is on the number of covariates included. As such, we will compare the proposed prior with the uniform and the Scott & Berger priors on the basis of the frequentist performances of posterior for the model size.

The paper is organised as follows. In Section 2 we define the notation used throughout the paper and formalise the problem of variable selection for linear regression models in a Bayesian framework. In Section 3 we discuss the current objective model priors for variable selection (the uniform prior and the Scott & Berger prior), and we also present the proposed prior based on losses. The results of simulation studies are provided in Section 4. In the section we examine the performance of the considered priors on the basis of the frequentist results of the corresponding model size posterior distributions. Section 5 reports the analysis results using two real data sets widely discussed in literature. The final Section 6 concludes and provides some discussion points on the proposed prior and its comparison with the other two priors.

2 Notation and problem specification

The Bayesian set up for variable selection problems is, as follows. The task is to explain a response variable by means of a set of d possible covariates, with a sample of observations of size n. Given the vector $\mathbf{y}$ of n responses, the design matrix $\mathbf{X}$ of size $n \times d$, an intercept α and a vector of coefficients $\boldsymbol{\beta}$ of dimension d, the response outcome y_i is expressed as

$$y_i = \alpha + \sum_{j=1}^{d} \beta_j X_{ij} + \varepsilon_i, \qquad i = 1, \ldots, n, \tag{1}$$

where $\varepsilon_i \sim N(0, 1/\phi)$, for $i = 1, \ldots, n$, are the i.i.d. normally distributed errors with unknown variance $1/\phi$. We assume that the number of observations n is larger than the number of covariates (i.e. $n > d$), and the design matrix is of full rank. The variable selection problem can be seen as identifying which of the possible d covariates has impact on $\mathbf{y}$. In other words, we aim to identify which of the regression parameters β_js are different from zero. Let us consider the binary vector $\boldsymbol{\gamma}$, where the j-th element is zero if $\beta_j = 0$ and one if $\beta_j \neq 0$. Then, the generic Bayesian regression model is indicated by

$$M_{\boldsymbol{\gamma}} = \left\{ f(\mathbf{y}|\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi); \pi_{\boldsymbol{\gamma}}(\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi) \right\}, \tag{2}$$

where

$$f(\mathbf{y}|\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi) = N(\mathbf{y}|\alpha + \mathbf{X}_{\boldsymbol{\gamma}}\boldsymbol{\beta}_{\boldsymbol{\gamma}}, \mathrm{I}/\phi),$$

and $\pi_{\boldsymbol{\gamma}}(\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi)$ represents the prior distribution for the parameters of the model, the so-called model-specific parameter prior. Note that the cardinality of $\boldsymbol{\gamma}$ gives the number of covariates included in model $M_{\boldsymbol{\gamma}}$. There are 2^d possible regression models, each one of them identified by $\boldsymbol{\gamma}$.

In the Bayesian framework, inference about model uncertainty is based on the model posterior probability

$$p(M_{\boldsymbol{\gamma}}|\mathbf{y}) \propto m(\mathbf{y}|M_{\boldsymbol{\gamma}}) p(M_{\boldsymbol{\gamma}}),$$

where $p(M_{\boldsymbol{\gamma}})$ is the prior probability for model $M_{\boldsymbol{\gamma}}$ and

$$m(\mathbf{y}|M_{\boldsymbol{\gamma}}) = \int f(\mathbf{y}|\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi) \pi_{\boldsymbol{\gamma}}(\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi)\, d\alpha\, d\boldsymbol{\beta}_{\boldsymbol{\gamma}}\, d\phi, \tag{3}$$

is the marginal likelihood of the observations under model $M_{\boldsymbol{\gamma}}$. The model posterior distribution can then be used to either estimate a specific regression model or perform model averaging.

The number of possible regression models grows exponentially with d. When d is large the posterior probabilities in (3) are often small for most of models, and posterior inclusion probabilities could give a better idea of the posterior uncertainty in comparison to model posterior probabilities. The posterior inclusion probability of the j-th covariate is defined as

$$\omega_j = \Pr(\gamma_j \neq 0|\mathbf{y}) = \sum_{\boldsymbol{\gamma}} p(M_{\boldsymbol{\gamma}}|\mathbf{y}) \cdot 1_{\gamma_j = 1}.$$

Prior and posterior inclusion probabilities originate from the common idea in Bayesian variable selection to consider variable inclusions as exchangeable Bernoulli trials, with the common inclusion probability ω, implying

$$p(M_{\boldsymbol{\gamma}}|\omega) = \omega^{|\boldsymbol{\gamma}|}(1-\omega)^{d-|\boldsymbol{\gamma}|}. \tag{4}$$

Model-specific parameter prior Prior choice on model-specific parameters has received much attention, and well-received priors include the Zellner–Siow prior (Zellner and Siow, 1980), the Zellner's g-prior (Zellner, 1986), the mixtures of g-priors (Liang *et al.*, 2008), and the more recent *robust* prior by Bayarri *et al.* (2012), among others. The robust prior has the advantage of yielding closed-form marginal likelihoods and to not suffer from the information paradox (Liang *et al.*, 2008). The robust prior is defined as

$$\begin{aligned} \pi_{\boldsymbol{\gamma}}(\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi) &= \pi(\alpha, \phi) \times \pi(\boldsymbol{\beta}_{\boldsymbol{\gamma}}) \\ &= \phi^{1/2} \int_0^{\infty} N_{|\boldsymbol{\gamma}|}(\boldsymbol{\beta}_{\boldsymbol{\gamma}} | 0, g\Sigma_{\boldsymbol{\gamma}}) \pi_{\boldsymbol{\gamma}}(g)\, dg, \end{aligned} \tag{5}$$

where $\Sigma_{\boldsymbol{\gamma}}$ is the covariance matrix of the maximum likelihood estimator of $\boldsymbol{\beta}_{\boldsymbol{\gamma}}$. The distribution of g is given by

$$\pi_{\boldsymbol{\gamma}}(g) = a[\rho_{\boldsymbol{\gamma}}(b+n)]^a (g+b)^{-(a+1)} \cdot 1_{\{g > \rho_{\boldsymbol{\gamma}}(b+n)-b\}},$$

where $a, b > 0$ and $\rho_{\boldsymbol{\gamma}} \geq b/(b+n)$. The prior (5) is called *robust* as its tails behave like the tails of a multivariate t density, therefore less sensible to outliers.

In this paper, as the focus is solely on model priors, every analysis is performed by using the same model-specific parameter prior, the robust prior with hyperparameter values $a = 1/2$, $b = 1$ and $\rho_{\boldsymbol{\gamma}} = 1/(d+1)$, as recommended in Bayarri *et al.* (2012), so differences in the results can be ascribed to differences in the model prior.

3 Model priors in objective variable selection

To the best of our knowledge, the choice of model prior probabilities which convey minimal information is limited to two options: the uniform prior and the so-called Scoot & Berger prior.

The model uniform prior is obtained by assigning equal prior mass to each regression model, that is $p(M_{\boldsymbol{\gamma}}) = 1/2^d$ for any $\boldsymbol{\gamma}$, and it yields a prior inclusion probability of $\omega = 1/2$. Scott and Berger (2010) discuss the following model prior for variable selection

$$\begin{aligned} p(M_{\boldsymbol{\gamma}}) &= \int_0^1 p(M_{\boldsymbol{\gamma}} | \omega) \pi_{\omega}(\omega)\, d\omega \\ &= \frac{1}{d+1} \binom{d}{|\gamma|}^{-1}. \end{aligned} \tag{6}$$

Their model prior is obtained by assigning a beta prior to ω, with both the hyper-parameters equal to one, and then marginalising over ω. This prior was previously discussed in the literature, see for example Ley and Steel (2009), with the aim of representing prior minimal information. The motivation behind the choice of the model prior by Scott and Berger (2010) lies in its property to correct for multiplicity, which can be seen as an issue when model choice is performed by multiple statistical testing with respect to a reference model (typically the null model or the full model). The Scott & Berger prior induces a marginal prior inclusion probability of $\omega = 1/2$ for each covariate, same as the uniform model prior. However, as thoroughly discussed in Scott and Berger (2010), given that their prior is function of d, it allows the multiplicity correction. It may also be worthwhile to mention that the choice of a uniform prior for ω, which is a probability of success of Bernoulli trials, somehow contradicts the usual choice of a Beta$(1/2, 1/2)$ as default prior in parameter inference (Jeffreys, 1961; Bernardo and Smith, 1994).

3.1 Model prior based on losses

The model prior based on losses has been introduced by Villa and Walker (2015). The basic idea is that we can assign a *worth* to each model by objectively measuring what is lost if the model is removed from the space of models, and it is the true one.

To derive the prior for $M_{\boldsymbol{\gamma}}$ we then follow the framework detailed in Villa and Walker (2015). First, we derive the loss in information when the true model is removed, which we indicate by $\mathrm{L}_I(M_{\boldsymbol{\gamma}})$. We apply the well-known Bayesian result in Berk (1966), which states that if a model is misspecified the posterior asymptotically accumulates on the nearest model, where the nearest model is the one which minimises the Kullback–Leibler divergence (Kullback and Leibler, 1951) from the true model. To illustrate, let us assume that we want to assess the *worth* of the regression model $M_{\boldsymbol{\gamma}}$, with the simplified notation $\boldsymbol{\theta}_{\boldsymbol{\gamma}} = (\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}}, \phi)$. The loss in information by removing model $M_{\boldsymbol{\gamma}}$, when it is the true model, is given by

$$L_I(M_{\boldsymbol{\gamma}}) = -\int_{\boldsymbol{\theta}_{\boldsymbol{\gamma}}} \inf_{\boldsymbol{\theta}_{\boldsymbol{\gamma}'} \neq \boldsymbol{\theta}_{\boldsymbol{\gamma}}} D_{KL}\Big(f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}}) \| f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'})\Big) \pi_{\boldsymbol{\gamma}}(\boldsymbol{\theta}_{\boldsymbol{\gamma}})\, d\boldsymbol{\theta}_{\boldsymbol{\gamma}},$$

where $f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'})$ represents the regression distribution of $M_{\boldsymbol{\gamma}'}$, that is, the regression model which is the most similar to $f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}})$. We then see that the loss in information associated to model $M_{\boldsymbol{\gamma}}$ is the expected minimum Kullback–Leibler divergence between $M_{\boldsymbol{\gamma}}$ and the nearest one, where the expectation is taken with respect to the prior $\pi_{\boldsymbol{\gamma}}(\boldsymbol{\theta}_{\boldsymbol{\gamma}})$, representing the prior uncertainty about the true values of α, $\boldsymbol{\beta}_{\boldsymbol{\gamma}}$ and ϕ.

To link the *worth* of a model to its prior probability, we use the *self-information* loss function, as discussed in Villa and Walker (2015). Briefly, this particular type of loss function (also known as the *log-loss* function) measures the performance of a probability statement with respect to an outcome. Thus, for every probability assignment $P = \{\Pr(A), A \in \Omega\}$, the self-information loss function is defined as

$$\mathrm{L}(P, A) = -\log \Pr(A).$$

More details and properties of this particular loss function can be found, for example, in Merhav and Feder (1998). As both the expected minimum Kullback–Leibler divergence between model $M_{\boldsymbol{\gamma}}$ and the nearest model, and the self-information loss function measure the same quantity, i.e. the loss in information, we can equate the two measurements yielding

$$p(M_{\boldsymbol{\gamma}}) \propto \exp\left\{\min_{\boldsymbol{\gamma}' \neq \boldsymbol{\gamma}} D_{KL}\Big(f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}}) \| f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'})\Big)\right\}. \tag{7}$$

However, as the number of covariates may easily be moderate to large, it is also necessary to take into considerations the complexity of the model. For the regression model $M_{\boldsymbol{\gamma}}$, we denote the loss due to complexity by $\mathrm{L}_C(M_{\boldsymbol{\gamma}})$, and it is determined as follows. If we keep model $M_{\boldsymbol{\gamma}}$ in the space of models, the loss would be proportional to the number of covariates that have to be considered and measured. Therefore, the loss of keeping a linear regression model increases with the number of covariates it contains, and we have

$$\mathrm{L}(\text{remove } M_{\boldsymbol{\gamma}}) = \mathrm{U}(\text{keep } M_{\boldsymbol{\gamma}}) = -c \cdot |\boldsymbol{\gamma}|,$$

and

$$\mathrm{L}(\text{keep } M_{\boldsymbol{\gamma}}) = c \cdot |\boldsymbol{\gamma}|, \qquad c > 0,$$

where $\mathrm{L}(\cdot)$ represents a loss and $\mathrm{U}(\cdot)$ a utility.

The loss component due to complexity is easily fit in our framework and the model prior for $M_{\boldsymbol{\gamma}}$ is

$$p(M_{\boldsymbol{\gamma}}) \propto \exp\left\{ \int_{\boldsymbol{\theta}_{\boldsymbol{\gamma}}} \left[\inf_{\boldsymbol{\gamma}' \neq \boldsymbol{\gamma}} D_{KL}\Big(f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}}) \| f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'}) \Big) \pi_{\boldsymbol{\gamma}}(\boldsymbol{\theta}_{\boldsymbol{\gamma}})\, d\boldsymbol{\theta}_{\boldsymbol{\gamma}} \right] - c \cdot |\boldsymbol{\gamma}| \right\}. \tag{8}$$

In other words, the prior is constructed based upon a cumulative loss with a component representing the loss in information and a component representing the loss due to complexity.

The following Theorem 1 (which proof is in the Appendix) shows the expression of the minimum Kullback–Leibler divergence between regression models.

Theorem 1 *Let* $M_{\boldsymbol{\gamma}} = \{ f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}}); \pi_{\boldsymbol{\gamma}}(\boldsymbol{\theta}_{\boldsymbol{\gamma}}) \}$ *and* $M_{\boldsymbol{\gamma}'} = \{ f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'}); \pi_{\boldsymbol{\gamma}'}(\boldsymbol{\theta}_{\boldsymbol{\gamma}'}) \}$ *be linear normal regression models as in* (2)*, with design matrices, respectively,* $\mathbf{X}_{\boldsymbol{\gamma}}$ *and* $\mathbf{X}_{\boldsymbol{\gamma}'}$*. If* $\mathbf{X}_{\boldsymbol{\gamma}}^T \mathbf{X}_{\boldsymbol{\gamma}'}$ *is invertible, the minimum Kullback–Leibler divergence between* $f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}})$ *and* $f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'})$ *is*

$$\min_{\boldsymbol{\theta}_{\boldsymbol{\gamma}'}} D_{KL}\Big(f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}}) \| f(\mathbf{y}|\boldsymbol{\theta}_{\boldsymbol{\gamma}'}) \Big) = 0 \qquad \forall\, \boldsymbol{\gamma} \neq \boldsymbol{\gamma}'. \tag{9}$$

Theorem 1 shows that the minimum Kullback–Leibler divergence between any two linear regression models is zero, regardless to the number of covariates in the models. This means that, in variable selection problems for linear regression models, there is no loss in information in selecting the "wrong" model, as such the model prior in (8) becomes

$$p(M_{\boldsymbol{\gamma}}) \propto \exp\left\{ -c \cdot |\boldsymbol{\gamma}| \right\}. \tag{10}$$

Figure 1 compares the proposed prior with the Scott & Berger prior, in terms of prior mass versus model size, for the case of $d = 30$ covariates. The uniform prior has been omitted from the plot as it would be represented by a horizontal line, reflecting a constant prior probability independent from the number of covariates included in the regression model. Whilst Scott & Berger prior has a symmetrical behaviour, the proposed prior assigns more mass to the more simple models than to the more complex ones, as expected from expression (10).

It is important to discuss some remarks about the proposed prior. First, the prior in (10) has practical relevance, especially in modern variable selection problems where the number of covariates can be significantly large. In fact, assigning less prior mass on large models reflects a parsimonious prior position, where the inclusion of additional covariates has to be supported by relatively strong evidence in the data.

The constant c plays an important role. As $c \to 0$, the prior (10) tends to be uniform, that is $p(M_{\boldsymbol{\gamma}}) \propto 1$. If, on the other hand, we increase the value of c, the proposed prior will put higher mass on small models, yielding more parsimonious regression models. A natural choice is to set $c = 1$, as if there is no prior information about the true model this appears to be the less informative choice. We empirically show in the simulation studies in Section 4 that $c = 1$ leads better results, when compared to the uniform and the Scott & Berger priors, if there is no prior information about the model size.

The proposed prior with the recommended value $c = 1$ does not correct the multiplicity. One way of solving this issue, if required, would be to set constant c as a function of the number of covariates. However, as mentioned in Section 1, we believe that for variable-selection problems the approach should be in line with the Bayesian framework of having prior and posterior probability representing, respectively, prior and posterior uncertainty. Although we agree that

multiplicity correction is one way to define model prior probabilities in an objective sense, we argue that this is not a necessary condition for a model prior to satisfy, in particular, if the problem of interest is related to prediction.

As a final remark, we note that the prior based on loses induces a binomial prior for model size k:

$$\begin{aligned} P(k) &= \binom{d}{k} e^{-ck} \Big/ \sum_{l=0}^{d} \binom{d}{l} e^{-cl} \\ &= \binom{d}{k} \left(\frac{1}{e^c+1}\right)^k \left(\frac{e^c}{e^c+1}\right)^{d-k}. \end{aligned} \tag{11}$$

Therefore, the marginal prior inclusion probability is given by

$$\omega = (e^c + 1)^{-1}.$$

The immediate consequence of the above result is that the prior inclusion probability depends on the value of the constant c. If we set, as recommended for a minimally informative prior, $c = 1$, then we obtain $\omega = 0.27$. A smaller values of c will lead a prior inclusion probability close to $1/2$, that is $\omega \to 1/2$ for $c \to 0$. On the other hand, as c increases, the prior inclusion probability becomes negligible and the model prior mass tends to be concentrated on the null model.

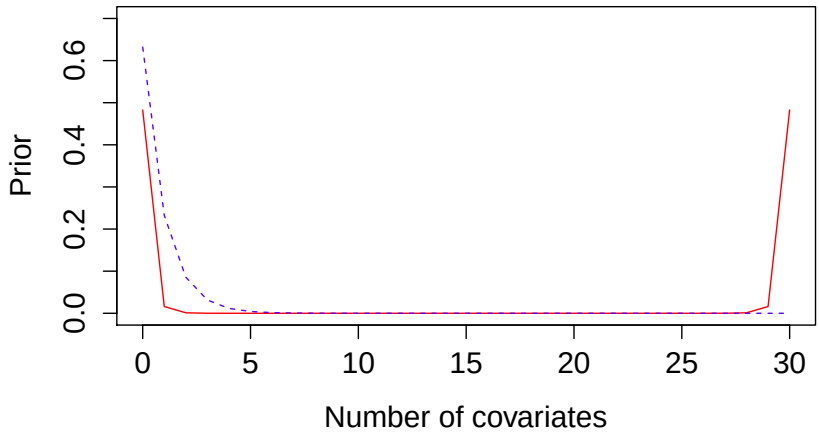

Figure 1: Prior model comparison when the full model contains 30 covariates: Scott & Berger prior (continuous red line) versus our proposed prior (dashed blue line).

4 Simulation study

In this section we show the results of a simulation study designed to compare the performance of the proposed model prior based on losses with the uniform prior and the Scott & Berger prior.

It is important to highlight here that the value of c considered for this simulation is one. The reason lies in the predictability of the behaviour of our prior for different values of c.

It is well known that a variable selection problem is driven by both the choice of the model prior and of the model-specific parameter prior. However, here the interest is in the effects on variable selection determined by the prior probability on the space of models. As such, the simulation exercise described has the purpose to analyse the frequentist properties of the posterior distribution on the model size and, to minimise any possible effects of the model-specific parameters prior, we choose to use the robust prior of Bayarri *et al.* (2012) in conjunction with the three model priors. We obtained very similar results using the Zellner–Siow prior for the parameter of the regression (results are not included in this paper).

In line with objective Bayesian analysis, we have examined two key frequentist properties of the posterior distribution for the model size: the coverage and the mean squared error. The general idea is to perform repeated simulations from similar regression models. For the coverage, we compute the proportion of times the true model size is included in the 95% credible interval of the posterior. We also compute the mean squared error from the mean and from the median of the true model size.

The simulation study encompasses a wide range of different scenarios, in terms of sample size, number covariates and model size, so to perform a thorough analysis. In detail, we have considered three sample sizes ($n = 30$, $n = 50$, and $n = 100$) and four maximum number of covariates ($d = 3$, $d = 5$, $d = 10$ and $d = 15$) so to be able to compare the priors for different ratios number of covariates to sample size. To avoid cumbersome and difficult to interpret results, we have generated the repeated samples from models with, relatively to d, a small, medium and large number of covariates. This has been accomplished by randomly selecting the number of covariates according to a fixed prior inclusion probability ω of 0.15, 0.50 and 0.75, respectively.

The above values of n, d and ω give a total of 36 cases. For each case we repeat the following procedure for 100,000 times:

- Generate a design matrix $\mathbf{X}$ of size $d \times n$ where each element is an independent realisation of a standard normal distribution;
- Generate a binary vector $\boldsymbol{\gamma}$ from a sequence of d independent Bernoulli experiments with probability of success equal to ω;
- From the robust prior in (5), we generate the vector of coefficients $\boldsymbol{\beta}_{\boldsymbol{\gamma}}$;
- We generate the response vector from the regression model in (1), considering $\phi = 1$;
- Finally, by using the above values of the design matrix and the vector of responses, we have computed the necessary quantities, including the marginal likelihoods, the model posteriors and the model size posterior distribution.

The last step of the above procedure has been performed under the three model priors considered in this paper: the uniform prior, the Scott & Berger prior and the proposed prior.

With our simulation experience, the true model was equally identified under all the three considered priors. This part of the study has been performed by computing the average number of false positive and false negative in detecting the coefficients, both considering the highest probability model (HPM) and the median probability model (MPM). By false positive we consider a covariate which was included in any of the above models when it was not present in the

true model. And by false negative we consider a covariate not in the true model but included in the MPM or the HPM. As the focus of this paper is on model prior and its performance on model size estimation the above results are not reported in detail. Instead, we discuss the effect of the model prior choice on the posterior model size distribution.

The detailed results of the simulations on the model size are presented in Table 1. Under each model prior, we report the coverage of the 95% credible interval of the posterior, the mean squared error from the mean and the mean squared error from the median. From the experiments, the following aspects emerge. The coverage under the Scott & Berger prior appears to be the most stable across sample size and model size, although it is constantly above the nominal value of 95%. The uniform prior shows a coverage too low with respect to the nominal value when $d > 5$ and $\omega = 0.15$. This is more obvious for sample sizes of $n = 30$ and $n = 50$. The prior we propose appears to have a relatively low coverage when $\omega = 0.75$. However, it shows an overall performance in the coverage closer to the nominal value than the Scott & Berger's.

To study the variability of posterior estimations, the mean squared errors from the posterior mean values and posterior median values are summarized in Table 1. As in terms of behaviour both mean squared errors have similar pattern, we examine in detail the mean squared error from the mean, which is also detailed in Figure 2. As one would expect, across all values of d and ω the mean squared errors decrease as n increases. If we focus on ω, we note the following. For $\omega = 0.15$, the uniform prior has the worse performance, while our proposed prior has generally the best performance. The only exception being the scenario with $n = 100$ and $d = 15$. The three priors have very similar performance for expected medium size regression models, i.e. $\omega = 0.50$. Finally, for relatively large models, our prior is outperformed by both the uniform prior and the Scott & Berger prior. However, in this latter case, the differences are not as prominent as the case of relatively small models and tend to disappear for large values of d.

Overall, it appears that the proposed prior with $c = 1$ has the most stable frequentist performance, making it appealing in scenarios of minimal prior knowledge about the true model size. On the other hand, if one had prior information about the true models size being close to the full model, a value of c smaller than one could be selected (e.g. $c = 0.5$) so to obtain a model prior more similar to the uniform prior and increase performance.

n = 30		**Uniform**			**Scott & Berger**			**Proposed,** $c = 1$		
d	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	ω
3	0.995	0.515	0.445	0.998	0.371	0.297	0.999	0.162	0.163	0.15
	0.995	0.382	0.507	0.998	0.421	0.605	0.967	0.515	0.685	0.50
	0.989	0.545	0.746	0.997	0.529	0.802	0.908	0.930	1.166	0.75
5	0.980	1.770	1.640	0.995	0.969	0.784	0.997	0.363	0.323	0.15
	0.992	0.652	0.763	0.997	0.793	1.059	0.950	0.886	1.163	0.50
	0.960	1.021	1.221	0.995	0.843	1.163	0.852	1.682	2.014	0.75
10	0.773	7.464	7.444	0.993	2.438	1.760	0.989	1.179	1.011	0.15
	0.990	0.398	0.475	0.991	0.421	0.480	0.966	0.401	0.493	0.50
	0.957	0.561	0.706	0.990	0.460	0.593	0.886	0.840	0.976	0.75
15	0.786	11.266	11.314	0.992	4.501	3.656	0.973	2.074	1.926	0.15
	0.996	0.107	0.129	0.996	0.102	0.117	0.974	0.113	0.130	0.50
	0.972	0.177	0.212	0.993	0.153	0.185	0.918	0.222	0.249	0.75

n = 50		**Uniform**			**Scott & Berger**			**Proposed,** $c = 1$		
d	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	ω
3	0.997	0.363	0.298	0.999	0.252	0.206	0.999	0.126	0.134	0.15
	0.995	0.312	0.417	0.998	0.347	0.478	0.971	0.399	0.524	0.50
	0.988	0.445	0.618	0.995	0.427	0.635	0.932	0.693	0.867	0.75
5	0.988	1.220	1.131	0.997	0.603	0.478	0.998	0.267	0.248	0.15
	0.992	0.546	0.665	0.995	0.678	0.888	0.954	0.719	0.956	0.50
	0.963	0.830	1.045	0.993	0.710	0.989	0.878	1.263	1.548	0.75
10	0.882	5.536	5.382	0.997	1.104	0.746	0.996	0.773	0.629	0.15
	0.991	0.378	0.456	0.993	0.404	0.463	0.970	0.383	0.482	0.50
	0.960	0.519	0.659	0.990	0.448	0.581	0.896	0.732	0.856	0.75
15	0.784	9.115	9.133	0.996	1.259	0.808	0.990	1.255	1.076	0.15
	0.995	0.115	0.133	0.996	0.113	0.128	0.974	0.119	0.138	0.50
	0.975	0.171	0.201	0.994	0.156	0.188	0.918	0.203	0.231	0.75

n = 100		**Uniform**			**Scott & Berger**			**Proposed,** $c = 1$		
d	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	Coverage	MSE Mean	MSE Median	ω
3	0.999	0.223	0.174	0.999	0.151	0.120	0.999	0.083	0.084	0.15
	0.995	0.245	0.323	0.995	0.277	0.364	0.978	0.298	0.382	0.50
	0.987	0.336	0.461	0.993	0.325	0.471	0.949	0.482	0.596	0.75
5	0.990	0.741	0.619	0.997	0.347	0.269	0.998	0.183	0.170	0.15
	0.993	0.418	0.532	0.992	0.531	0.661	0.967	0.499	0.658	0.50
	0.967	0.613	0.810	0.994	0.537	0.764	0.909	0.851	1.051	0.75
10	0.954	3.270	2.989	0.998	0.420	0.279	0.997	0.449	0.340	0.15
	0.991	0.372	0.441	0.993	0.415	0.476	0.970	0.375	0.471	0.50
	0.958	0.508	0.647	0.990	0.451	0.580	0.894	0.701	0.821	0.75
15	0.866	5.810	5.716	0.997	0.269	0.152	0.994	0.660	0.530	0.15
	0.996	0.102	0.120	0.996	0.108	0.123	0.977	0.109	0.125	0.50
	0.975	0.160	0.187	0.994	0.149	0.178	0.922	0.187	0.210	0.75

Table 1: Summary of simulation results for the three priors: uniform prior, Scott & Berger prior and the proposed prior with $c = 1$. A proportion of identifying the true number of covariates within the 95% of the posterior probability (Coverage) and the mean squared error of posterior mean values (MSE Mean) and posterior median values (MSE Median) are computed.

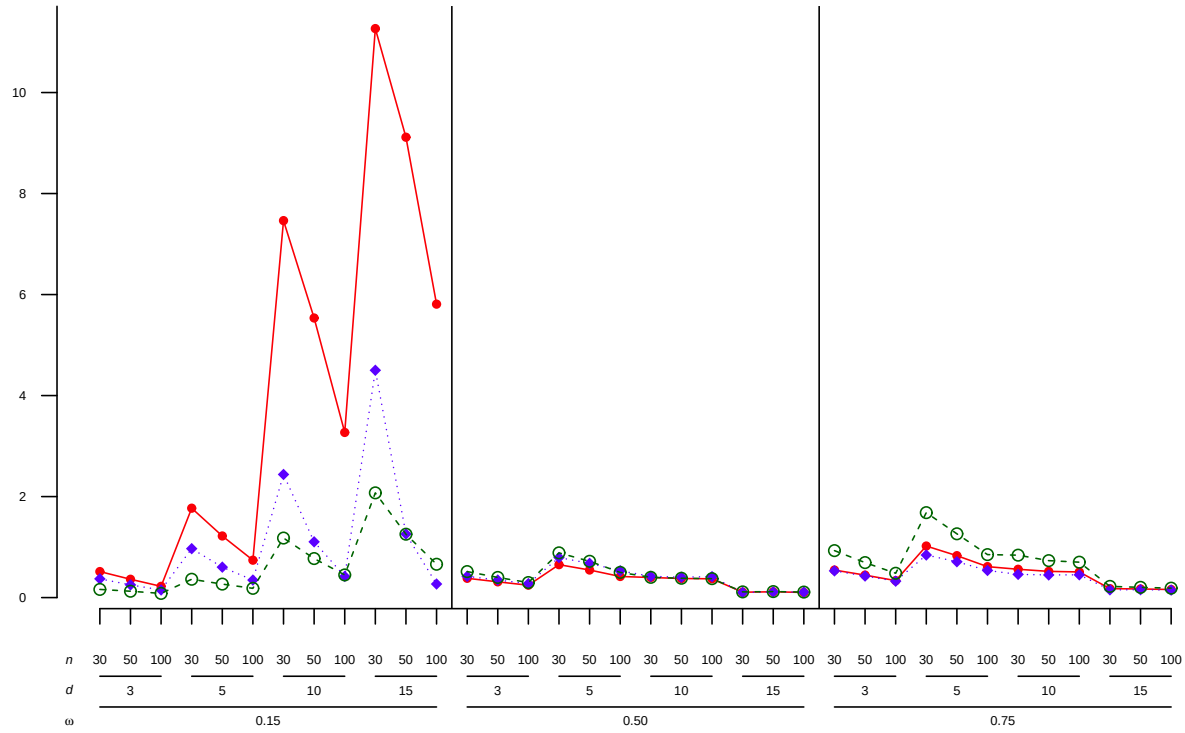

Figure 2: Mean squared error from the mean from the simulation exercise on the posterior for the model size. The plot reports the results obtained by applying on the model space the uniform prior (red continuous line), Scott & Berger prior (blue dotted line) and the proposed prior (dashed green line).

5 Illustrative examples with real data sets

In this section we investigate the properties of objective model priors for variable selection in real data sets. The considered data sets are US crime data (Vandaele, 1978) and the Hald data (Woods *et al.*, 1932), which have been extensively used in literature (see Kubinyi (1996) and Liang *et al.* (2008), for example). At last, we perform a robustness analysis to highlight the sensitivity of the inferential results to relatively small changes to the data.

5.1 US crime data

The first data set consists of $n = 47$ observations for $d = 15$ covariates related to crime data in the US for 1960 (Vandaele, 1978). The aim of the analysis is to understand which variables, such as population type, police expenditure and ethnicity (among others), impact on the crime rate per head in the population (the response variable).

The summary statistics of the model size posterior distributions are presented in Table 2, where we considered the proposed prior with different values of the constant c. The plot of the posterior distributions for the model size are represented in Figure 3. The mean and the median obtained by the use of the uniform prior and the Scott & Berger prior are similar. The proposed prior, with $c = 1$, gives mean and median model size smaller than the ones obtained by using the other two priors. We note that the posterior mean and median model size increase when c is smaller than one, as expected, and for $c > 1$ these values decrease as the prior puts more mass on smaller models. Again, this result is in line on what discussed in Sections 3 and 4. From the table, and by inspecting the the posterior plots, we see that the Scott & Berger prior yields a posterior for the model size with higher uncertainty than the one obtained by using either the uniform prior or the proposed prior.

The HPM identified by the uniform prior includes 6 covariates, unlike the one determined by Scott & Berger prior and the proposed prior, which both assign the highest mass to a model of size 3. The posterior probability for the HPM is of 0.019 for the uniform prior, 0.021 for the Scott & Berger prior, and 0.056 when the proposed prior (with $c = 1$) is used. In Table 3 we can see what covariates are included in the HPM under each one of the considered priors, and that both Scott & Berger and our prior agree in the result. In Table 3, the proposed model prior tends to assign lower probabilities compared to the other two priors, particularly for covariates which are unlikely to be included (posterior inclusion probabilities less than 1/2). The main result is that the covariates concerning age of males, unemployment rate among male and probability of imprisonment are included in the MPM under the uniform and the Scott & Berger priors, but not under the proposed prior.

Although the proposed prior and Scott & Berger's disagree on the size of the MPM, we note that both values, 3 and 6, are well within the 95% credible interval of the respective posteriors. The same can not be said for the uniform prior, as the corresponding credible interval does not contain the size 3.

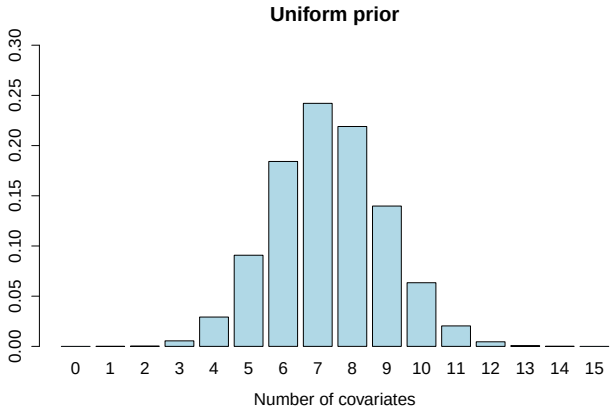

(a)

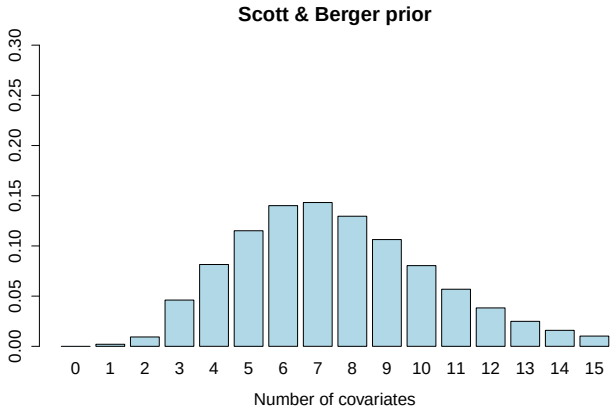

(b)

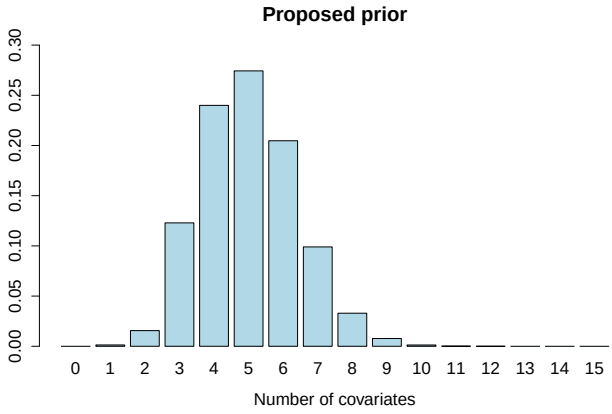

(c)

Figure 3: Posterior distribution of the number of covariates for the US Crime data set. Three model priors are considered: uniform, Scott & Berger and proposed prior with $c = 1$.

Model prior	Mean	Median	SD	95% C.I.	HPM	MPM
Uniform	7.32	7	1.78	(4,11)	6	6
Scott & Berger	7.49	7	2.91	(3,14)	3	6
Proposal($c = 1.0$)	5.00	5	1.56	(3,8)	3	3
Proposal($c = 0.5$)	6.06	6	1.72	(4,10)	3	3
Proposal($c = 1.5$)	4.13	4	1.42	(2,7)	3	3
Proposal($c = 2.0$)	3.45	3	1.29	(1,6)	3	3

Table 2: Comparison of the posterior summary statistics for the US Crime data set. Four statistics for the number of covariates are measured: mean, median, standard deviation (SD) and the 95% confidence interval (95% C.I.). The number of covariates included in the highest posterior probability model (HPM) and the median probability model (MPM) are reported.

	Posterior Inclusion Probability		
Covariate	Uniform	Scott & Berger	Proposed ($c = 1$)
Percentage of males aged 14-24	**0.76•**	**0.70**	0.46
Indicator variable for a Southern state	0.25	0.28	0.10
Mean years of schooling	**0.88•**	**0.85•**	**0.72•**
Police expenditure in 1960	**0.83•**	**0.84•**	**0.81•**
Police expenditure in 1959	0.38	0.41	0.27
Labour force participation rate	0.23	0.27	0.11
Number of males per 1000 females	0.38	0.42	0.29
State population	0.28	0.31	0.12
Number of non-whites per 1000 people	0.23	0.27	0.10
Unemployment rate of urban males 14-24	0.30	0.34	0.11
Unemployment rate of urban males 35-39	**0.52•**	**0.51**	0.21
Gross domestic product per head	0.36	0.37	0.18
Income inequality	**0.98•**	**0.97•**	**0.96•**
Probability of imprisonment	**0.70•**	**0.66**	0.44
Average time served in state prisons	0.25	0.28	0.12

Table 3: Posterior inclusion probabilities for the US Crime data set. The covariates with posterior inclusion probabilities greater than 1/2 are highlighted in bold, and a dot notation represents the covariate included in the highest posterior probability model.

5.2 Hald data

The Hald data set contains $n = 13$ observations with $d = 4$ covariates, and it concerns an engineering application to study the cement composition (Woods *et al.*, 1932). In particular, the study considers the effect on the heat evolved per gram of cement (in calories) by the amount of tricalcium aluminate, the amount of tricalcium silicate, the amount of tricalcium aluminio ferrite and the amount of dicalcium silicate.

The variable selection problem for the Hald data turns out to be a much simpler task than the one discussed in Section 5.1. In fact, by examining both Table 4 and Figure 4, we note that the size of HPM and MPM are 2 under each considered prior. The only exception being the proposed prior with $c = 2$. In this case, as expected, the prior is extremely parsimonious assigning high mass on the null model, in comparison to the mass assigned to any other regression model. There is a slightly higher variability in the posterior of the model size under the Scott & Berger prior, and this is reflected by both a higher standard deviation and a larger credible interval. The coinciding result on the model size estimation is supported by the similar posterior inclusion probabilities in Table 5. The posterior inclusion probability assigned to the first two covariates is definitely above 1/2 under each prior, whilst the one assigned to the remaining two covariates is clearly below 1/2.

The Hald data set appears to represent a case where the information contained in the data about which covariate should be included is so strong that the choice of the model prior does not make much difference; provided, of course, that this choice is not too extreme. The above is reflected in the posterior probability associated to the HPM being close or above 50% under each prior. In fact, the probability is 0.55 under the uniform prior, 0.47 under the Scott & Berger prior and 0.67 under the proposed prior.

	Mean	Median	SD	95% C.I.	HPM	MPM
Uniform	2.27	2	0.74	(2,3)	2	2
Scott & Berger	2.41	2	0.81	(2,4)	2	2
Proposal ($c = 1$)	2.12	2	0.72	(2,3)	2	2
Proposal ($c = 0.5$)	2.25	2	0.73	(2,3)	2	2
Proposal ($c = 1.5$)	2.07	2	0.71	(2,3)	2	2
Proposal ($c = 2$)	0.02	0	1.15	(0,0)	2	2

Table 4: Comparison of the posterior summary statistics for the Hald data set. Four statistics for the number covariates are measured: mean, median, standard deviation (SD) and the 95% confidence interval (95% C.I.). The number of covariates included in the highest posterior probability model (HPM) median probability model (MPM) are reported.

5.3 Robustness of model priors

To analyse the robustness of the proposed prior to small changes in the data, and to compare the results with the uniform and the Scott & Berger priors, we look at the performance in estimating the model size and the posterior inclusion probabilities for repeated random sub-samples of the original data, both for the US Crime data and the Hald data.

The general idea is to extract 500 sub-samples of a size that is the 85% of the numerosity of the data set, and perform variable selection as done in Section 5 using these sub-samples as ob-

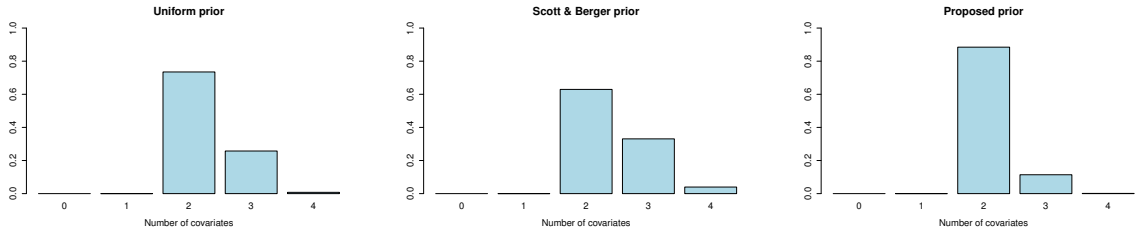

(a) (b) (c)

Figure 4: Posterior distribution of the number of covariates for the Hald data set. Three model priors are considered: uniform, Scott & Berger and proposed prior with $c = 1$.

	Posterior Inclusion Probability		
Covariate	Uniform	Scott & Berger	Proposed ($c = 1$)
Ticalcium aluminate	**0.98•**	**0.98•**	**0.99•**
Tricalcium silicate	**0.75•**	**0.76•**	**0.75•**
Tetracalcium alumino ferrite	0.18	0.26	0.08
Dicalcium silicate	0.36	0.42	0.30

Table 5: Posterior inclusion probabilities for the Hald data set. The covariates with a posterior inclusion probability greater than 1/2 are highlighted in bold, and a dot notation represents the covariate included in the highest posterior probability model.

servations. Then we compare the performance of the model priors by considering the histogram of the posterior mean model size, and the box-plots of the posterior inclusion probabilities for the considered priors.

For the US Crime data set analysed in Section 5.1, we consider 500 samples of size 40. The analysis takes also into consideration the results for different values of c for the proposed prior. Figure 5 shows the histogram of the mean model size when we apply the uniform prior and the Scott & Berger; for the proposed prior, we consider $c = 1.0$, $c = 0.5$, $c = 1.5$ and $c = 2.0$. In the plots we have indicated with a vertical line the estimated model size when the whole data set is considered. We note the following. First, the uncertainty in the model size is larger when the Scott & Berger prior is used. All the remaining cases appear to lead similar variability in the posterior model size estimates. Second, there is a sensitivity of our prior to c, with an impact on the mean location of the distribution, in the sense that when c increases the mean model size decreases. This agrees with what discussed in Section 3 about the parsimony of the prior when c increases. Also, note that proposed prior tends to yield a posterior similar to the posterior yielded by the uniform prior for $c = 0.5$, as expected.

For what it concerns the posterior inclusion probabilities, in Figure 6, we see a relatively small impact of c on the variability of the probabilities, except when $c = 2.0$ when we see a larger variability for covariate number 13, for example. Effects of different values of the constant c appear in the mean of the estimates. For example, covariates 1 and 3. When we compare our prior with $c = 1$, to the uniform prior and, in particular, Scott & Berger's, it is interesting to focus on the three covariates that are included in the MPM under these two last priors but not under our prior. In Figure 6 these covariates are indicated by numbers 1, 11 and 14. We note that the size of the variability of their posterior inclusion probabilities under the proposed prior

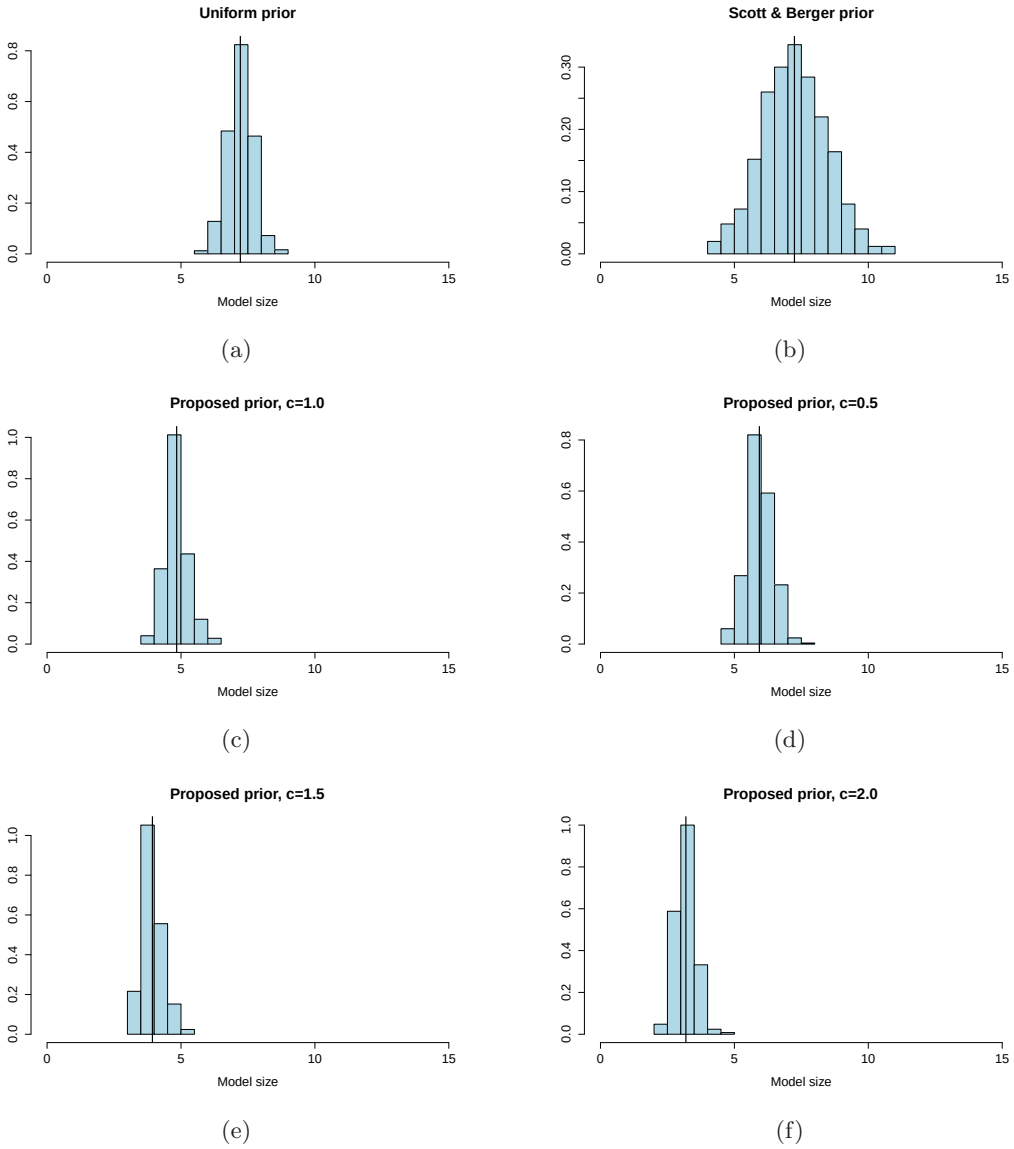

Figure 5: Histograms of the posterior mean model size for the considered priors, where the proposed prior has been analysed for $c = 1.0$, $c = 0.5$, $c = 1.5$ and $c = 2.0$ (500 sub-samples of size 40 from the US Crime data set).

admits possible inclusion of the covariates in the MPM. This seems to suggest that, although our prior has undoubtedly a feature of estimating regression models with less covariates than the one defined by the other two priors, it possibly reflects a low level of certainty in the data on whether the covariates have to be considered or not. Considering the covariates included in the MPM, the covariates 3, 4 and 13 in Figure 6, we can make the following comments. For the covariates 4 and 13, there is little to argue as the distributions of the respective posterior inclusion probabilities show strong information in the data in favour of their inclusion in the MPM model. However, with similar arguments for the above non included covariates, the variability of the posterior

inclusion probability of covariate 3 expresses a less strong support from the observations on whether the covariate has to be part of the regression model.

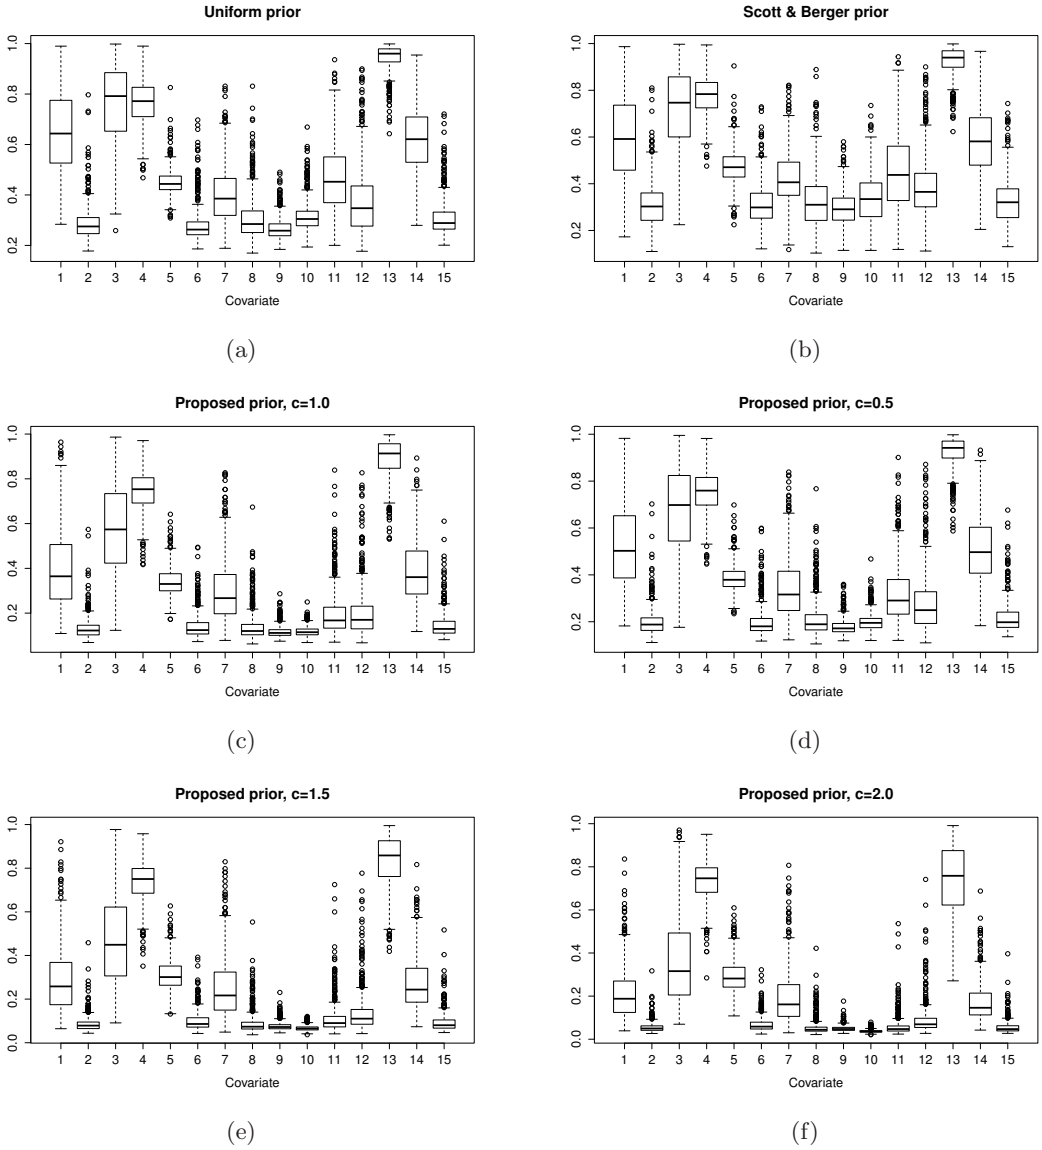

Figure 6: Box-plots of the posterior inclusion probabilities for the considered priors, where the proposed prior has been analysed for $c = 1.0$, $c = 0.5$, $c = 1.5$ and $c = 2.0$. Results are based on 500 replicates of sub-samples size of 40 from the US Crime data set.

The robustness analysis of the Hald data set, with a similar procedure as above, consisted in 500 repeated samples of size 10. The histograms of the posterior mean model size are in Figure 7, while the box-plots of the posterior inclusion probabilities are in Figure 8. As already discussed in Section 5.2, the information about the covariates that should be part of the regression model contained in the data appears to be strong, and this is reflected in the robustness of the approach. In particular, the histograms of the posterior mean model size are concentrated around the mean

obtained by considering the whole data set (vertical line). The inspection of the box-plots of

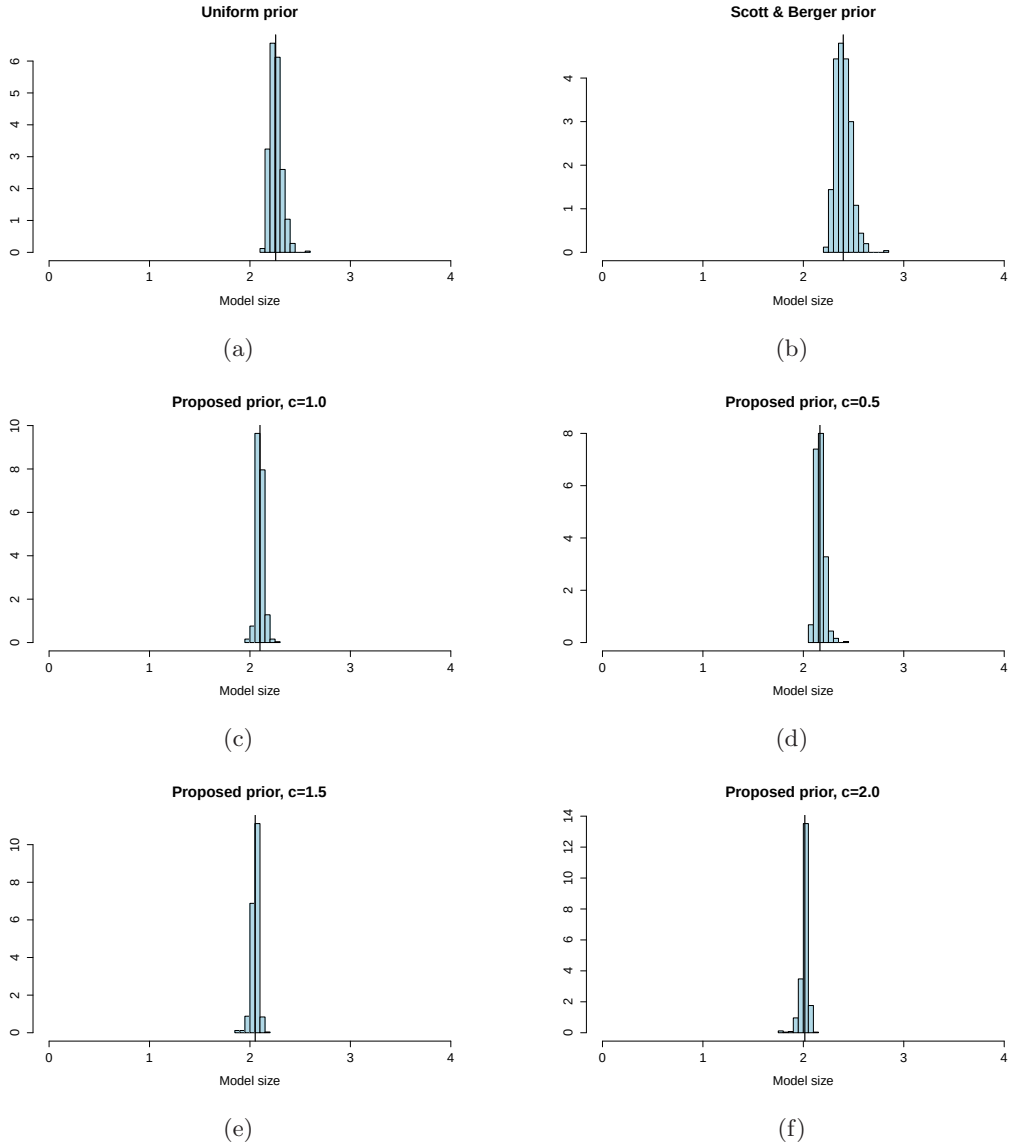

Figure 7: Histograms of the posterior mean model size for the considered priors, where the proposed prior has been analysed for $c = 1.0$, $c = 0.5$, $c = 1.5$ and $c = 2.0$ (500 sub-samples of size 10 from the Hald data set).

the posterior inclusion probabilities for this data set is straightforward. The impact of c is more noticeable on the covariates with either a high posterior inclusion probability (i.e. number 1) or with a low posterior inclusion probability (i.e. number 3). In fact, for c increasing, these probabilities become more prominent. In general, however, it doesn't seem that different priors lead to significantly different variability or shape of the distribution of the posterior inclusion probabilities.

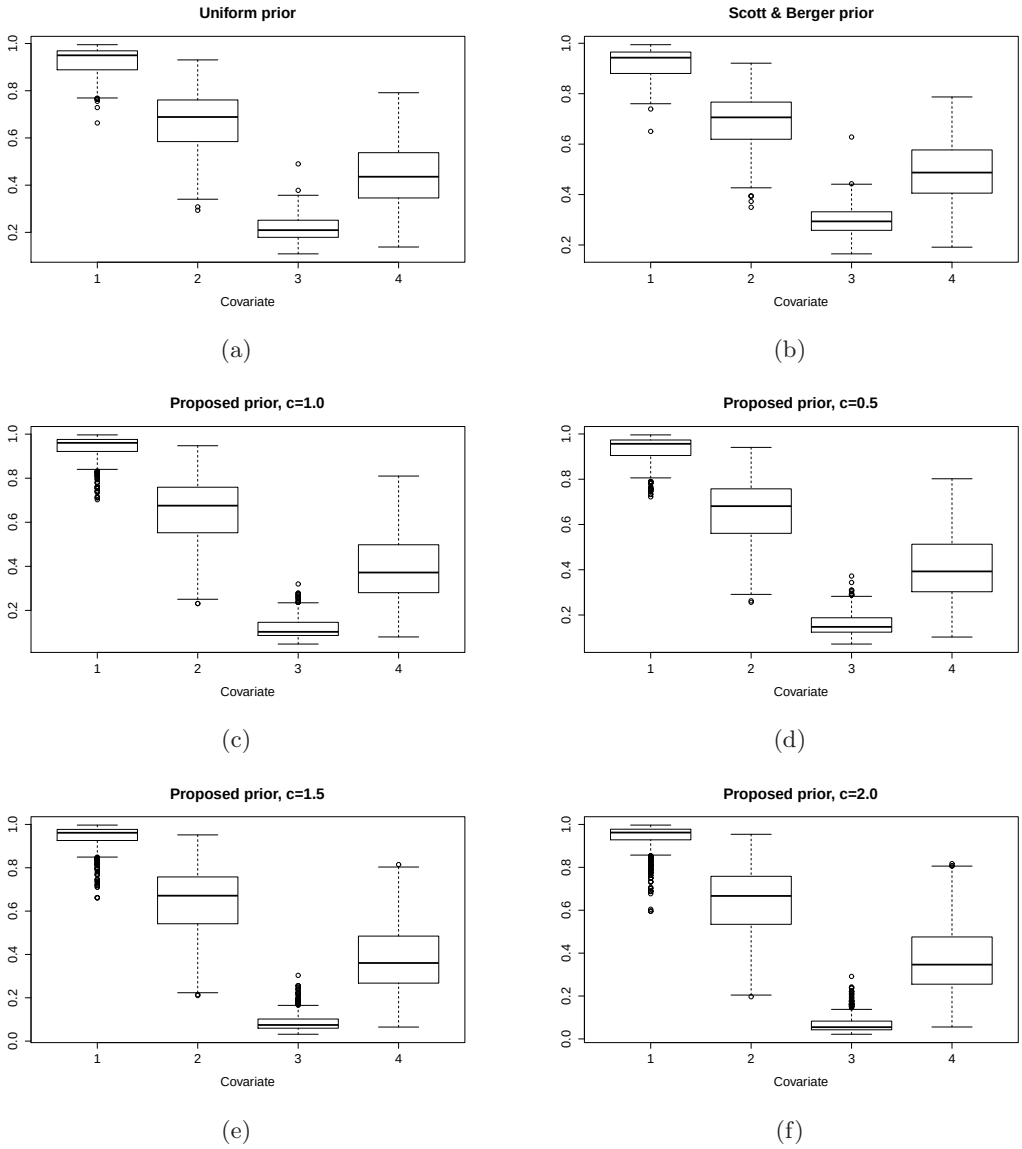

Figure 8: Box-plots of the posterior inclusion probabilities for the considered priors, where the proposed prior has been analysed for $c = 1.0$, $c = 0.5$, $c = 1.5$ and $c = 2.0$. Results are based on 500 replicates of sub-samples size of 10 from the Hald data set.

6 Discussion

This paper introduces a novel prior distribution for the model space in variable selection for linear regression. The prior is based on the idea that, if the "wrong" model is chosen, we incur in a loss that has two components: one represents the loss in information objectively measured by considering the *worth* of the model, and one related to the complexity of the model expressed by its size. The obtained model has the following two appealing properties. First, it has a simple expression. Second, it has the flexibility of representing minimal prior information about the model as well as to allow the inclusion of appropriate prior information with regard to the

model size.

The simulations studies carried out show good frequentist performance of the model prior when compared to other two existing minimally informative priors: the uniform prior and the Scott & Berger prior (Scott and Berger, 2010). In particular, when we set the constant $c = 1$, which is our recommendation for a prior carrying minimal prior information, the proposed prior outperforms the other two priors when the size of the true model is relatively small, with respect to the full model. In cases where the true model is close to the full model our prior is not performing as well as the other two priors. However, the difference in the performance is small and, if prior information is available, this lack of performance can be compensated by an appropriate choice of the constant c reflecting the prior knowledge.

When it comes to real data analysis, we note a fair closeness of the results obtained by using the proposed prior with the ones of Scott & Berger's. Both of them are more parsimonious than the uniform prior, unless the data itself is highly informative about which covariates should be included in the regression model. It is however important to highlight that the observed discrepancies between the proposed prior and the Scott & Berger prior are limited to the median probability model. In fact, the estimated highest probability models are the same for both the analysed data sets, and it is noteworthy to mention that the posterior probability associated to the highest probability model under our prior is the largest.

Appendix - Proof of Theorem 1

Let $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$ be non-identical binary vectors. With the corresponding design matrices $\mathbf{X}_{\boldsymbol{\gamma}}$ and $\mathbf{X}_{\boldsymbol{\gamma}'}$ and the vectors of coefficients $\boldsymbol{\beta}_{\boldsymbol{\gamma}}$ and $\boldsymbol{\beta}_{\boldsymbol{\gamma}'}$, the models are defined as

$$M_{\boldsymbol{\gamma}} : \mathbf{y} \sim N_n\left(\mathbf{1}_n\alpha + \mathbf{X}_{\boldsymbol{\gamma}}\boldsymbol{\beta}_{\boldsymbol{\gamma}}, \mathrm{I}/\phi\right) = N_n\left(\tilde{\mathbf{X}}_i\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}, \mathrm{I}/\phi\right)$$

$$M_{\boldsymbol{\gamma}'} : \mathbf{y} \sim N_n\left(\mathbf{1}_n\alpha + \mathbf{X}_{\boldsymbol{\gamma}'}\boldsymbol{\beta}_{\boldsymbol{\gamma}'}, \mathrm{I}/\phi\right) = N_n\left(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}, \mathrm{I}/\phi\right)$$

where $\tilde{\mathbf{X}}_{\boldsymbol{\gamma}} = [\mathbf{1}, \mathbf{X}_{\boldsymbol{\gamma}}]$, $\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'} = [\mathbf{1}, \mathbf{X}_{\boldsymbol{\gamma}'}]$, $\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}} = (\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}})^T$, and $\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'} = (\alpha, \boldsymbol{\beta}_{\boldsymbol{\gamma}'})^T$. The Kullback–Leibler divergence of $M_{\boldsymbol{\gamma}}$ from $M_{\boldsymbol{\gamma}'}$ is

$$\begin{aligned}
& D_{KL}(M_{\boldsymbol{\gamma}} \| M_{\boldsymbol{\gamma}'}) \\
= & \int m(\mathbf{y}|M_{\boldsymbol{\gamma}}) \Bigg\{ -\frac{\log(1/\phi)}{2} - \frac{\phi}{2}(\mathbf{y} - \tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}})^T(\mathbf{y} - \tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}) + \frac{\log(1/\phi)}{2} \\
& + \frac{\phi}{2}(\mathbf{y} - \tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T(\mathbf{y} - \tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}) \Bigg\} d\mathbf{y} \\
= & \int m(\mathbf{y}|M_{\boldsymbol{\gamma}})\phi \Bigg\{ -\frac{1}{2}\mathbf{y}^T\mathbf{y} + (\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}})^T\mathbf{y} - \frac{1}{2}(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}})^T(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}) + \frac{1}{2}\mathbf{y}^T\mathbf{y} - (\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T\mathbf{y} \\
& + \frac{1}{2}(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}) \Bigg\} d\mathbf{y} \\
= & \frac{\phi}{2}\Big((\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}})^T(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}) - 2(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}) + (\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}) \Big)
\end{aligned}$$

The first derivative of $D_{KL}(M_{\boldsymbol{\gamma}}\|M_{\boldsymbol{\gamma}'})$ with respect to $\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}$ is

$$\frac{\partial D_{KL}(M_{\boldsymbol{\gamma}}\|M_{\boldsymbol{\gamma}'})}{\partial \tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}} = \phi\Big(-(\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}})^T + (\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'})^T\Big)\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}.$$

If $\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}^T$ is invertible, $D_{KL}(M_{\boldsymbol{\gamma}}\|M_{\boldsymbol{\gamma}'})$ is minimized when $\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'} = (\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}^T\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'})^{-1}\tilde{\mathbf{X}}_{\boldsymbol{\gamma}'}^T\tilde{\mathbf{X}}_{\boldsymbol{\gamma}}\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}}$ and the minimum is given by

$$\min_{\tilde{\boldsymbol{\beta}}_{\boldsymbol{\gamma}'}} D_{KL}(M_{\boldsymbol{\gamma}}\|M_{\boldsymbol{\gamma}'}) = 0.$$

References

Barbieri, M. and Berger, J. O. (2004). Optimal predictive model selection. *Ann. Statist.* **32**, 870–897

Bayarri, M. J., Berger, J. O., Forte, A. and García-Donato, G. (2012). Criteria for Bayesian model choice with application to variable selection. *J. Ann. Statist.* **40**, 1550–1577

Bogdan, M., Ghosh J. and Tokar, S. T. (2008). Selecting explanatory variables with the modified version of the Bayesian information criterion. *Quality and Reliability Engineering International* **24**, 627–641

Berger, J. O. and Molina, G. (2005). Posterior model probabilities via path-based pairwise priors. *Statistica Neerlandica* **59**, 3–15

Berk, R. H. (1966). Limiting behaviour of posterior distributions when the model is incorrect. *Ann. of Math. Statist.* **37**, 51–58

Bernardo, J. M. and Smith, A. F. M. (1994). Bayesian Theory. New York: John Wiley & Sons.

Brown, P. J., Vannucci, M., and Fearn, T. (1998). Bayesian wavelength selection in multi-component analysis. *J. of Chemometrics* **12**, 173–182

Carlin, B. and Louis, T. (2000). Empirical Bayes: Past, present and future. *J. Amer. Statist. Assoc.* **95**, 1286–1289

Casella, G. and Moreno, E. (2006). Objective Bayesian variable selection. *J. Amer. Statist. Assoc.* **101**, 157–167

Clyde, M. A. and George, E. I. (2004). Model uncertainty. *Statistical Science* **19**, 81–94

Cui, W. and George, E. I. (2008). Empirical Bayes vs. fully Bayes variable selection. *J. Statist. Plann. Inference.* **138**, 888–900

Fernández, C., Ley, E. and Steel, M. F. J. (2001). Benchmark priors for Bayesian model averaging. *J. of Econometrics* **100**, 381–427

Garćia-Donato, G. and Martínez-Beneito, M. A. (2013). On sampling strategies in Bayesian variable selection problems with large model spaces. *J. Amer. Statist. Assoc.* **108**, 340–352

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A. and Rubin, D. B. (2004). *Bayesian Data Analysis*. Chapman and Hall/CRC

George, E. I. and Foster, D. P. (2000). Calibration and empirical Bayes variable selection. *Biometrika*. **87**, 731–747

Hoeting, J. A., Madigan, D., Raftery, A. E. and Volinsky, C. T. (1999). Bayesian model averaging: a tutorial. *Statist. Science* **14**, 382–401

Jeffreys, H. (1961). *Theory of Probability*. London: Oxford University Press.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *J. Amer. Statist. Assoc.* **90**, 773–795

Bubinyi, H. (1996). Evolutionary variable selection in regression and PLS analyses. *Chemometrics* **10**, 119–133

Kullback, S. and Leibler, R.A. (1951). On information and sufficiency. *Ann. of Math. Statist.* **22**, 79–86

Ley, E. and Steel, M. F. (2009). On the effect of prior assumptions in Bayesian model averaging with applications to growth regression. *J. Appl. Econometrics.* **24**, 651–674

Liang, F., Paulo, R., Molina, G., Clyde, M. and Berger, J. O. (2008). Mixtures of g-priors for Bayesian variable selection. *J. Amer. Statist. Assoc.* **103**, 410–423

Merhav, N. and Feder, M. (1998). Universal prediction. *IEEE Trans. Inf. Theory* **44**, 2124–2147

Nott, D. J. and Kohn, R. (2005). Adaptive sampling for Bayesian variable selection. *Biometrika* **92**, 747–763

O'Hara, R. B. and Sillanpää, M. J. (2009). A Review of Bayesian Variable Selection Methods: What, How and Which. *Bayesian Analysis* **4**, 85–118

Raftery, A. E., Madigan, D. and Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *J. Amer. Statist. Assoc.* **92**, 179–191

Shively, T. S., Kohn, R. and Wood, S. (1999). Variable selection and function estimation in additive nonparametric regression using a data-based prior. *J. Amer. Statist. Assoc.* **447**, 777–794

Scott, J. C. and Berger, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in variable-selection problems. *Ann. Statist.* **38**, 2587–2619

Vandaele, W. (1978). Participation in illegitimate activities: Ehrlich revisited. In *Deterrence and Incapacitation*. US National Academy of Sciences, 270–335.

Villa, C. and Walker, S. G. (2015). An objective Bayesian criterion to determine model prior probabilities. *Scandinavian Journal of Statistics* **42**, 947–966

Woods, H., Steinour, H. and Starke, H. (1932). Effect of Composition of Porland Cement on Heat Evolved During Hardening. *Industrial and Engineering Chemistry Research* **24**, 1207–1214

Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In *Bayesian inference and Decision Techniques: Essays in Honour of Bruno de Finetti*, Goel PK, Zellner A (eds). North-Holland: Amsterdam, 233–243

Zellner, A. and Siow, A. (1980). Posterior odds ratios for selected regression hypotheses. (with discussion). In *Bayesian Statistics*, Bernardo JM, DeGroot MH, Lindley DV, Smith AFM (eds). University Press: Valencia, 585–603